



Oracle

1Z0-117 Exam

Oracle Database 11g Release 2: SQL Tuning

Thank you for Downloading 1Z0-117 exam PDF Demo

You can Buy Latest 1Z0-117 Full Version Download

<https://www.certkillers.net/Exam/1Z0-117>

<https://www.certkillers.net>

Question: 1

Examine the query and its execution plan:

```
SQL > SELECT cust_last_name, sum (nu12(o.customer_id, 0, 1)) "Count"
      FROM customer c, orders o
     WHERE c.credit_limit > 1000
       AND c.customer_id = o.customer_id(+)
    GROUP By cust_last_name;
```

ID	Operations	Name	Rows	Bytes	Cost (%CPU)
0	SELECT STATEMENT		168	3192	6 (17)
1	HASH GROUP BY		168	3192	6 (17)
* 2	NESTED LOOPS OUTER		260	4940	5 (0)
* 3	TABLE ACCESS FULL	CUSTOMER	260	3900	5 (0)
* 4	INDEX RANGE SCAN	ORD_CUSTOMER_IX	105	420	0 (0)

Predicate Information (identified by operation id):

```
3 - filter ("C"."CREDIT_LIMIT">> 1000)
4 - access("C"."CUSTOMER_ID" = "O" "CUSTOMER_ID" (+))
   Filter ("O" "CUSTOMER_ID" (+)>0)
```

Which statement is true regarding the execution plan?

- A. This query first fetches rows from the CUSTOMERS table that satisfy the conditions, and then the join return NULL from the CUSTOMER_ID column when it does not find any corresponding rows in the ORDERS table.
- B. The query fetches rows from CUSTOMERS and ORDERS table simultaneously, and filters the rows that satisfy the conditions from the resultset.
- C. The query first fetches rows from the ORDERS table that satisfy the conditions, and then the join returns NULL form CUSTOMER_ID column when it does not find any corresponding rows in the CUSTOMERS table.
- D. The query first joins rows from the CUSTOMERS and ORDERS tables and returns NULL for the ORDERS table columns when it does not find any corresponding rows in the ORDERS table, and then fetches the rows that satisfy the conditions from the result set.

Answer: A

Question: 2

Which three statements are true about histograms?

- A. They capture the distribution of different values in an index for better selectivity estimates.
- B. They can be used only with indexed columns.
- C. They provide metadata about distribution of and occurrences of values in a table column.
- D. They provide improved selectivity estimates in the presence of data skew, resulting in execution

plans with uniform distribution.

E. They help the optimizer in deciding whether to use an index or a full table scan.

F. They help the optimizer to determine the fastest table join order.

Answer: C, E, F

Explanation:

C: A histogram is a frequency distribution (metadata) that describes the distribution of data values within a table.

E: It's well established that histograms are very useful for helping the optimizer choose between a full-scan and an index-scan.

F: Histograms may help the Oracle optimizer in deciding whether to use an index vs. a full-table scan (where index values are skewed) or help the optimizer determine the fastest table join order. For determining the best table join order, the WHERE clause of the query can be inspected along with the execution plan for the original query. If the cardinality of the table is too-high, then histograms on the most selective column in the WHERE clause will tip-off the optimizer and change the table join order.

Note:

* The Oracle Query Optimizer uses histograms to predict better query plans. The ANALYZE command or DBMS_STATS package can be used to compute these histograms.

Incorrect:

B: Histograms are NOT just for indexed columns.

– Adding a histogram to an un-indexed column that is used in a where clause can improve performance.

D: Histograms Opportunities

Any column used in a where clause with skewed data

Columns that are not queried all the time

Reduced overhead for insert, update, delete

Question: 3

```
SQL > EXPLAIN PLAN SET Statement_id = 'test' for
      SELECT prod_category, avg(amount_sold)
        FROM sales s, products p
       WHERE p.prod_id = s.prod_id
      GROUP BY prod_Category;
```

Explained.

```
SQL> SELECT id "id", parent_id, position "pos"
      lpad(' ', 2 level) || operations || decode(id, 0, 'cost=' || POSITION) "operations"
    Options "option" object_name "object"
  FROM plan_table
 Connect by prior id_parent_id START WITH id = 0
 ORDER BY id;
```

id	PARENT_ID	POS	Operation	Option	Object
0		539	SELECT STATEMENT Cost = 539		
1	0	1	HASH	Group By	
2	1	1	HASH JOIN		
3	2	1	VIEW		
4	3	1	HASH	GROUP BY	
5	4	1	PARTITION RANGE	ALL	
6	5	1	TABLE ACCESS	FULL	SALES index\$_joins\$_002
7	2	2	VIEW		
8	7	1	VIEW RANGE		
9	8	1	INDEX	FAST FULL SCAN	PRODUCTS_PK
10	8	2	INDEX	FAST FULL SCAN	PRODUCTS_PROD_CAT_IX

11 rows are selected

View the exhibit and examine the query and its execution plan from the PLAN_TABLE. Which statement is true about the execution?

- A. The row with the ID column having the value 0 is the first step execution plan.
- B. Rows are fetched from the indexes on the PRODUCTS table and from the SALES table using full table scan simultaneously, and then hashed into memory.
- C. Rows are fetched from the SALES table, and then a hash join operator joins with rows fetched from indexes on the PRODUCTS table.
- D. All the partitions of the SALES table are read in parallel.

Answer: C

Question: 4

Which four statements are correct about communication between parallel execution process?

- A. The number of logical pathways between parallel execution producers and consumers depends on the degree parallelism.
- B. The shared pool can be used for parallel execution messages buffers.
- C. The large pool can be used for parallel execution messages buffers.
- D. The buffer cache can be used for parallel execution message buffers.
- E. Communication between parallel execution processes is never required if a query uses full partition-wise joins.
- F. Each parallel execution process has an additional connection to the parallel execution coordinator.

Answer: A, B, E, F

Explanation:

A: Note that the degree of parallelism applies directly only to intra-operation parallelism. If inter-operation parallelism is possible, the total number of parallel execution servers for a statement can be twice the specified degree of parallelism. No more than two sets of parallel execution servers can run simultaneously. Each set of parallel execution servers may process multiple operations. Only two sets of parallel execution servers need to be active to guarantee optimal inter-operation parallelism.

B: By default, Oracle allocates parallel execution buffers from the shared pool.

F: When executing a parallel operation, the parallel execution coordinator obtains parallel execution servers from the pool and assigns them to the operation. If necessary, Oracle can create additional parallel execution servers for the operation. These parallel execution servers remain with the operation throughout job execution, then become available for other operations. After the statement has been processed completely, the parallel execution servers return to the pool.

Reference: Oracle Database Data Warehousing Guide, Using Parallel Execution

Question: 5

You have enabled DML by issuing: ALTER session ENABLE PARALLEL DML;

The PARELLEL_DEGREE_POLICY initialization parameter is set to AUTO.

Which two options true about DML statements for which parallel execution is requested?

- A. Statements for which PDML is requested will execute serially estimated time is less than the time specified by the PARALLEL_MIN_THRESHOLD parameter.
- B. Statements for which PDML is requested will be queued if the number of busy parallel execution servers greater than PARALLEL_MIN_SERVERS parameter.
- C. Statements for which PDML is requested will always execute in parallel if estimated execution in parallel if estimated execution time is greater than the time specified by the PARELLEL_MIN_TIME_THRESHOLD parameter.
- D. Statements for which PDML is requested will be queued if the number of busy parallel execution servers is greater than PARELLEL_SERVERS_TARGET parameter.
- E. Statement for which PDML is requested will be queued if the number of busy parallel execution servers is greater than PARELLEL_DEGREE_LIMIT parameter.

Answer: C, D

Explanation:

C: PARALLEL_MIN_TIME_THRESHOLD specifies the minimum execution time a statement should have before the statement is considered for automatic degree of parallelism. By default, this is set to 30 seconds. Automatic degree of parallelism is only enabled if PARALLEL_DEGREE_POLICY is set to AUTO or LIMITED.

D: PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle will queue SQL statements that require parallel execution, if the necessary parallel server processes are not available. Statement queuing will begin once the number of parallel server processes active on the system is equal to or greater than PARALLEL_SERVER_TARGET.

Note:

* PARALLEL_DEGREE_POLICY specifies whether or not automatic degree of Parallelism, statement

queuing, and in-memory parallel execution will be enabled.

AUTO

Enables automatic degree of parallelism, statement queuing, and in-memory parallel execution.

* PARALLEL_MIN_SERVERS specifies the minimum number of parallel execution processes for the instance. This value is the number of parallel execution processes Oracle creates when the instance is started.

Reference: Oracle Database Reference; PARALLEL_SERVERS_TARGET

Question: 6

Examine Exhibit1 to view the query and its AUTOTRACE output.

```
SQL> SET AUTOTRACE TRACEONLY
SQL> SELECT prod_category, AVG(amount_sold)
      FROM sales s,products P
     WHERE P.prod_id = S.prod_id
   GROUP BY prod_category;
```

Execution Plan

Plan hash volume: 1197568639							
id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time	Pstart
0	SELECT STATEMENT		5	255	639 (11)	00:00:07	
1	HASH GROUP BY		5	255	539 (11)	00:00:07	
2	HASH JOIN		72	3672	538 (11)	00:00:07	
3	VIEW	VW_GBC_5	72	2160	535 (11)	00:00:07	
4	HASH GROUP BY		72	648	535 (11)	00:00:07	
5	PARTITION RANGE ALL		918K	8075K	494 (3)	00:00:06	1
6	TABLE ACCESS FULL	SALES	918K	8075K	684 (3)	00:00:06	1
7	VIEW	INDEX\$_JOIN_002	72	1512	3 (34)	00:00:01	
8	HASH JOIN						
9	INDEX FAST FULL SCAN	PRODUCT_PK	72	1512	1 (0)	00:00:01	
10	INDEX FAST FULL SCAN	PRODUCT_PROD_CAT_IK	72	1512	1 (0)	00:00:01	

Predicate information (identified by operation id)"

```
2-access ("P". "PROD_ID" = "ITEM_1")
8-access (ROEID=ROWID)
```

Statistics

0	recursive calls
0	db block gets
1726	consistent gets
0	physical reads
0	Redo size
778	bytes sent via SQL "Net to client"
434	bytes received via SQL "Net from client"
2	SQL "Net roundtrips to/from client"
0	sorts (memory)

Which two statements are true about tracing?

- A. The displayed plan will be stored in PLAN_TABLE.
- B. Subsequent execution of this statement will use the displayed plan that is stored in v\$SQL.
- C. The displayed plan may not necessarily be used by the optimizer.
- D. The query will not fetch any rows; it will display only the execution plan and statistics.
- E. The execution plan generated can be viewed from v\$SQLAREA.

Answer: A, C

Explanation:

A: The PLAN_TABLE is automatically created as a public synonym to a global temporary table. This temporary table holds the output of EXPLAIN PLAN statements for all users. PLAN_TABLE is the default sample output table into which the EXPLAIN PLAN statement inserts rows describing execution plans

Incorrect:

B: V\$SQL lists statistics on shared SQL area without the GROUP BY clause and contains one row for each child of the original SQL text entered. Statistics displayed in V\$SQL are normally updated at the end of query execution. However, for long running queries, they are updated every 5 seconds. This makes it easy to see the impact of long running SQL statements while they are still in progress.

D: autotrace traceonly – Displays execution plan and statistics without displaying the returned rows. This option should be used when a large result set is expected.

E: V\$SQLAREA lists statistics on shared SQL area and contains one row per SQL string. It provides statistics on SQL statements that are in memory, parsed, and ready for execution.

Note:

* The autotrace provides instantaneous feedback including the returned rows, execution plan, and statistics. The user doesn't need to be concerned about trace file locations and formatting since the output is displayed instantly on the screen. This is very important data that can be used to tune the SQL statement.

* SET AUTOTRACE ON

The AUTOTRACE report includes both the optimizer execution path and the SQL statement execution statistics.

SET AUTOTRACE TRACEONLY

Similar to SET AUTOTRACE ON, but suppresses the printing of the user's query output, if any. If STATISTICS is enabled, query data is still fetched, but not printed.

Question: 7

Which two types of column filtering may benefit from partition pruning?

- A. Equally operates on range-partitioned tables.
- B. In-list operators on system-partitioned tables
- C. Equality operators on system-partitioned tables
- D. Operators on range-partitioned tables
- E. Greater than operators on hash-partitioned tables

Answer: A, D

Explanation:

The query optimizer can perform pruning whenever a WHERE condition can be reduced to either one of the following two cases:

partition_column = constant

partition_column IN (constant1, constant2, ..., constantN)

In the first case, the optimizer simply evaluates the partitioning expression for the value given, determines which partition contains that value, and scans only this partition. In many cases, the equal sign can be replaced with another arithmetic comparison, including <, >, <=, >=, and <>. Some queries using BETWEEN in the WHERE clause can also take advantage of partition pruning.

Note:

* The core concept behind partition pruning is relatively simple, and can be described as “Do not scan partitions where there can be no matching values”.

When the optimizer can make use of partition pruning in performing a query, execution of the query can be an order of magnitude faster than the same query against a nonpartitioned table containing the same column definitions and data.

* Example:

Suppose that you have a partitioned table t1 defined by this statement:

```
CREATE TABLE t1 (
    fname VARCHAR(50) NOT NULL,
    lname VARCHAR(50) NOT NULL,
    region_code TINYINT UNSIGNED NOT NULL,
    dob DATE NOT NULL
)
PARTITION BY RANGE( region_code ) (
    PARTITION p0 VALUES LESS THAN (64),
    PARTITION p1 VALUES LESS THAN (128),
    PARTITION p2 VALUES LESS THAN (192),
    PARTITION p3 VALUES LESS THAN MAXVALUE
);
```

Consider the case where you wish to obtain results from a query such as this one:

```
SELECT fname, lname, region_code, dob
  FROM t1
 WHERE region_code > 125 AND region_code < 130;
```

It is easy to see that none of the rows which ought to be returned will be in either of the partitions **p0** or **p3**; that is, we need to search only in partitions **p1** and **p2** to find matching rows. By doing so, it is possible to expend much less time and effort in finding matching rows than would be as *pruning*.

Question: 8

Which two statements about In-Memory Parallel Execution are true?

- A. It can be configured using the Database Resource Manager.
- B. It increases the number of duplicate block images in the global buffer cache.
- C. It requires setting PARALLEL_DEGREE_POLICY to LIMITED.
- D. Objects selected for In-Memory Parallel Execution have blocks mapped to specific RAC instances.
- E. It requires setting PARALLEL_DEGREE_POLICY to AUTO
- F. Objects selected for In-Memory Parallel Execution must be partitioned tables or indexes.

Answer: D, E

Explanation:

D, E: In-Memory Parallel Execution

When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database decides if an object that is accessed using parallel execution would benefit from being cached in the SGA (also called the buffer cache). The decision to cache an object is based on a well-defined set of heuristics including the size of the object and frequency on which it is accessed. In an Oracle RAC environment, Oracle Database maps pieces of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle Database automatically knows which buffer cache to access to find different parts or pieces of the object. Using this information, Oracle Database prevents multiple instances from reading the same information from disk over and over again, thus maximizing the amount of memory that can cache objects. If the size of the object is larger than the size of the buffer cache (single instance) or the size of the buffer cache multiplied by the number of active instances in an Oracle RAC cluster, then the object is read using direct-path reads.

E: PARALLEL_DEGREE_POLICY specifies whether or not automatic degree of Parallelism, statement queuing, and in-memory parallel execution will be enabled.

AUTO

Enables automatic degree of parallelism, statement queuing, and in-memory parallel execution.

Incorrect:

C:

LIMITED

Enables automatic degree of parallelism for some statements but statement queuing and in-memory Parallel Execution are disabled. Automatic degree of parallelism is only applied to those statements that access tables or indexes decorated explicitly with the PARALLEL clause. Tables and indexes that have a degree of parallelism specified will use that degree of parallelism.

Reference: Oracle Database VLDB and Partitioning Guide 11g, How Parallel Execution Works

Question: 9

Which two are benefits of In-Memory Parallel Execution?

- A. Reduction in the duplication of block images across multiple buffer caches
- B. Reduction in CPU utilization
- C. Reduction in the number of blocks accessed
- D. Reduction in physical I/O for parallel queries
- E. Ability to exploit parallel execution servers on remote instance

Answer: A, C

Explanation:

Note: In-Memory Parallel Execution

When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database decides if an object that is accessed using parallel execution would benefit from being cached in the SGA (also called the buffer cache). The decision to cache an object is based on a well-defined set of heuristics including the size of the object and frequency on which it is accessed. In an Oracle RAC environment, Oracle Database maps pieces of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle Database automatically knows which buffer cache to access to find different parts or pieces of the object. Using this information, Oracle Database prevents multiple instances from reading the same information from disk over and over again, thus maximizing the amount of memory that can cache objects. If the size of the object is larger than the size of the buffer cache (single instance) or the size of the buffer cache multiplied by the number of active instances in an

Oracle RAC cluster, then the object is read using direct-path reads.

Reference: Oracle Database VLDB and Partitioning Guide 11g, How Parallel Execution Works

Question: 10

You plan to bulk load data `INSERT INTO . . . SELECT FROM` statements. Which two situations benefit from parallel `INSERT` operations on tables that have no materialized views defined on them?

- A. Direct path insert of a million rows into a partitioned, index-organized table containing one million rows and a conventional B*tree secondary index.
- B. Direct path insert of a million rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index.
- C. Direct path insert of 10 rows into a partitioned, index-organized table containing one million rows and conventional B* tree secondary index.
- D. Direct path insert of 10 rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index
- E. Conventional path insert of a million rows into a nonpartitioned, heap-organized containing 10 rows and having a conventional B* tree index.
- F. Conventional path insert of 10 rows into a nonpartitioned, heap-organized table one million rows and a bitmapped index.

Answer: A, B

Explanation:

Note:

- * A materialized view is a database object that contains the results of a query.
- * You can use the `INSERT` statement to insert data into a table, partition, or view in two ways: conventional `INSERT` and direct-path `INSERT`.
- * With direct-path `INSERT`, the database appends the inserted data after existing data in the table. Data is written directly into datafiles, bypassing the buffer cache. Free space in the existing data is not reused. This alternative enhances performance during insert operations and is similar to the functionality of the Oracle direct-path loader utility, SQL*Loader. When you insert into a table that has been created in parallel mode, direct-path `INSERT` is the default.
- * Direct-path `INSERT` is not supported for an index-organized table (IOT) if it is not partitioned, if it has a mapping table, or if it is reference by a materialized view.
- * When you issue a conventional `INSERT` statement, Oracle Database reuses free space in the table into which you are inserting and maintains referential integrity constraints
- * Conventional `INSERT` always generates maximal redo and undo for changes to both data and metadata, regardless of the logging setting of the table and the archivelog and force logging settings of the database

Question: 11

Which are the two prerequisites for enabling star transformation on queries?

- A. The `STAR_TRANSFORMATION_ENABLED` parameter should be set to TRUE or TEMP_DISABLE.
- B. A B-tree index should be built on each of the foreign key columns of the fact table(s),

- C. A bitmap index should be built on each of the primary key columns of the fact table(s).
- D. A bitmap index should be built on each of the foreign key columns of the fact table(s).
- E. A bitmap index must exist on all the columns that are used in the filter predicates of the query.

Answer: A, E

Explanation:

A: Enabling the transformation

E: Star transformation is essentially about adding subquery predicates corresponding to the constraint dimensions. These subquery predicates are referred to as bitmap semi-join predicates. The transformation is performed when there are indexes on the fact join columns (s.timeid, s.custid...). By driving bitmap AND and OR operations (bitmaps can be from bitmap indexes or generated from regular B-Tree indexes) of the key values supplied by the subqueries, only the relevant rows from the fact table need to be retrieved. If the filters on the dimension tables filter out a lot of data, this can be much more efficient than a full table scan on the fact table. After the relevant rows have been retrieved from the fact table, they may need to be joined back to the dimension tables, using the original predicates. In some cases, the join back can be eliminated.

Star transformation is controlled by the star_transformation_enabled parameter. The parameter takes 3 values.

TRUE - The Oracle optimizer performs transformation by identifying fact and constraint dimension tables automatically. This is done in a cost-based manner, i.e. the transformation is performed only if the cost of the transformed plan is lower than the non-transformed plan. Also the optimizer will attempt temporary table transformation automatically whenever materialization improves performance.

FALSE - The transformation is not tried.

TEMP_DISABLE - This value has similar behavior as TRUE except that temporary table transformation is not tried.

The default value of the parameter is FALSE. You have to change the parameter value and create indexes on the joining columns of the fact table to take advantage of this transformation.

Reference: Optimizer Transformations: Star Transformation

Question: 12

An application accessing your database got the following error in response to SQL query:

ORA-12827: insufficient parallel query slaves available

View the parallel parameters for your instance:

NAME	TYPE	VALUE
fast_start_parallel_rollback	string	LOW
parallel_adaptive_multi_tuning	boolean	TRUE
parallel_automatic_turing	boolean	FALSE
parallel_degree_limit	string	32
parallel_degree_policy	string	LIMITED
parallel_execution_message_size	integer	16384
parallel_force_local	boolean	FALSE
parallel_io_cap_enabled	boolean	FALSE
parallel_max_servers	integer	128
parallel_min_servers	integer	50
parallel_server	Integer	0
parallel_server_instances	string	AUTO
parallel_server	boolean	1
parallel_servers_target	integer	8
parallel_threads_servers_per_cpu	integer	2

No hints are used and the session use default parallel settings. What four changes could you make to help avoid the error and ensure that the query executes in parallel?

- A. Set PARELLEL_DEGREE_POLICY to AUTO.
- B. Increase the value of PARELLEL_MAX_SERVERS.
- C. Increase PARELLEL_SERVERS_TARGET.
- D. Decrease PARELLEL_MIN_PERCENT.
- E. Increase PARELLEL_MIN_SERVERS.
- F. Decrease PARELLEL_MIN_TIME_THRESHOLD.
- G. Increase PARELLEL__MIN_TIME_THRESHOLD.

Answer: A, C, D, G

Explanation:

C: PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle will queue SQL statements that require parallel execution, if the necessary parallel server processes are not available. Statement queuing will begin once the number of parallel server processes active on the system is equal to or greater than PARALLEL_SERVER_TARGET.

By default, PARALLEL_SERVER_TARGET is set lower than the maximum number of parallel server processes allowed on the system (PARALLEL_MAX_SERVERS) to ensure each parallel statement will get all of the parallel server resources required and to prevent overloading the system with parallel server processes.

D:

Note: ORA-12827: insufficient parallel query slaves available

Cause: PARALLEL_MIN_PERCENT parameter was specified and fewer than minimum slaves were acquired

Action: either re-execute query with lower PARALLEL_MIN_PERCENT or wait until some running queries are completed, thus freeing up slaves

A, G: PARALLEL_MIN_TIME_THRESHOLD specifies the minimum execution time a statement should

have before the statement is considered for automatic degree of parallelism. By default, this is set to 30 seconds. Automatic degree of parallelism is only enabled if PARALLEL_DEGREE_POLICY is set to AUTO or LIMITED.

Thank You for trying 1Z0-117 PDF Demo

To Buy Latest 1Z0-117 Full Version Download visit link below

<https://www.certkillers.net/Exam/1Z0-117>

Start Your 1Z0-117 Preparation

[Limited Time Offer] Use Coupon “CKNET” for Further discount on your purchase. Test your 1Z0-117 preparation with actual exam questions.

<https://www.certkillers.net>